

CLAIMS

What is claimed is:

1. A system comprising:
a first node that includes an ordering point for data, the first node being operative to employ a write-back transaction associated with writing the data back to memory, the first node broadcasting a write-back message to at least one other node in the system in response to an acknowledgement provided by the memory indicating that the ordering point for the data has migrated from the first node to the memory.
2. The system of claim 1, wherein the first node comprises a processor having an associated cache that comprises a plurality of cache lines, one of the plurality of cache lines having an associated state that defines the cache line as a cache ordering point for the data prior to employing the write-back transaction.
3. The system of claim 1, wherein the at least one other node provides a response to the first node acknowledging receipt of the write-back message broadcast by the first node.
4. The system of claim 3, wherein the first node maintains the write-back transaction active until the first node receives responses from the at least one other node to the write-back message broadcast by the first node.
5. The system of claim 4, further comprising a third node that issues a source broadcast request for the data employing a source broadcast protocol, the third node retrying the source broadcast request for the data in response to recognizing a conflict associated with the source broadcast request for the data.
6. The system of claim 5, wherein the conflict is recognized by the third node in response to one of (i) receiving the write-back message broadcast by the first node while the source-broadcast request for the data is active at the third node, or (ii) receiving a conflict response from the first node to the source broadcast request issued by the third node.

7. The system of claim 5, wherein the third node retries the source broadcast request employing a forward progress protocol.
8. The system of claim 1, wherein the first node further comprises a request engine having an associated miss address file, the request engine allocating an entry in the miss address file associated with the write-back transaction for the data that is maintained in the miss address file until responses have been received from all other nodes in the system to the write-back message broadcast by the first node.
9. A computer system, comprising:
 - a first processor that provides a write-back request to transfer an ordering point for desired data from associated cache of the first processor to memory;
 - the memory providing an acknowledgement back to the first processor in response to the write-back request, the first processor providing a source broadcast write-back request to the system in response to the acknowledgement provided by the memory; and
 - at least one other processor in the system that provides an acknowledgement response to the first processor in response to the source broadcast write-back request provided by the first processor.
10. The system of claim 9, wherein the system employs a source broadcast protocol, the system further comprising a third node that issues a source broadcast request for the desired data, the third node reissuing the request in response to recognizing a conflict associated with the source broadcast request for the desired data.
11. The system of claim 10, wherein the conflict is recognized by the third node in response to one of (i) receiving the source broadcast write-back request provided by the first node while the source-broadcast request for the desired data is active at the third node, or (ii) receiving a conflict response from the first node to the source broadcast request issued by the third node.
12. The system of claim 9, wherein the third node reissues the request employing a forward progress protocol implemented in the system.

13. The system of claim 9, further comprising an entry in a miss address file at the first processor that is associated with transferring the ordering point from the associated cache of the first processor to the memory, the entry in the miss address file being maintained until responses to the source broadcast write-back request have been received from all other processors in the system.

14. The system of claim 9, wherein the first processor comprises a cache line that contains the desired data in a state that defines the cache line as the ordering point for the desired data prior to issuing the write-back request to the memory.

15. The system of claim 14, wherein the state that defines the cache line as the ordering point for the desired data is selected from a group consisting of a modified state, an owner state and a dirty state, the cache line transitioning to an invalid state after issuing the write-back request to the memory.

16. A multiprocessor computer system, comprising:
means for issuing a write-back request to migrate an ordering point for data from an associated cache to memory; and
means for providing a source broadcast write-back message associated with the data to the system in response to the memory acknowledging receipt of the write-back request.

17. The system of claim 16, further comprising means at each of at least one node in the system for acknowledging receipt of the source broadcast write-back message.

18. The system of claim 16, further comprising means for retiring an outstanding transaction associated with migration of the ordering point to the memory from the associated cache in response to receiving the acknowledgement of receipt of source broadcast write-back message.

19. The system of claim 16, further comprising means for recognizing a conflict associated with the data.

20. The system of claim 16, wherein the means for issuing the write-back request comprises a first processor including a cache line that contains the data in a state that defines the cache line as the ordering point for the data prior to migration of the ordering point to the memory.

21. The system of claim 20, wherein the state that defines the cache line as the ordering point for the data is selected from a group consisting of a modified state, an owner state and a dirty state, the cache line transitioning to an invalid state after issuing the write-back request to the memory.

22. The system of claim 16, further comprising means for transitioning a cache line in the associated cache to an invalid cache state for the data after the write-back request is issued.

23. A method comprising:

providing a write-back request from a first processor node to transfer an ordering point associated with data to memory;
acknowledging receipt of the write back request at the memory; and
issuing a source broadcast message from the first processor node to other nodes in response to acknowledging receipt of the write back request at the memory.

24. The method of claim 23, further comprising providing a response from each of the other nodes to acknowledge receipt of the source broadcast message at the other nodes.

25. The method of claim 24, further comprising maintaining a write-back transaction at the first processor node until the first processor node receives the responses to the source broadcast message from the other nodes.

26. The method of claim 25, further comprising retiring the write-back transaction at the first processor node in response to receiving the responses from each of the other nodes acknowledging receipt of the source broadcast message.

27. The method of claim 23, further comprising recognizing a conflict associated with a request for the data provided by at least one of the other nodes.

28. The method of claim 27, wherein the request for the data provided by the at least one of the other nodes comprises a source broadcast request, the recognizing of the conflict further comprising recognizing the conflict in response to one of (i) receiving the source broadcast write-back request provided by the first processor node while the source-broadcast request for the data is outstanding at the at least one of the other nodes, or (ii) receiving a conflict response from the first node to a source broadcast request issued by the at least one of the other nodes.

29. The method of claim 23, wherein the first processor comprises a cache line that contains the data in a state that defines the cache line as the ordering point for the data prior to issuing the write-back request to the memory.

30. The method of claim 29, further comprising transitioning the state of the data in the first processor from a first state to an invalid state after issuing the write-back request to the memory, the first state being selected from a group consisting of a modified state, an owner state and a dirty state.